

Geant4 Liquid Argon Validation

Isaac Harris

Supervisor: Hans Wenzel

July 29, 2016

This paper summarizes my work as an intern at Fermi National Accelerator Laboratory over the summer. I worked in multiple different areas, so the paper does not flow naturally from section to section; the transitions are rather rough. Nonetheless, a lot of the work is concerned with validating Geant4 processes, especially in liquid Argon time projection chambers. This may be of interest to many of the current and future experiments that use liquid argon TPCs, ie. LArIAT, DUNE, etc. The specific process we looked at were hadronic cross sections of pions, kaons, and nucleons in liquid Argon, the shape of electromagnetic showers, and separation of muons and pions by charge in the absence of an electric and magnetic field.

We were also interested in how detectors are simulated by programs such as LArSoft and LArG4. After the event is generated by Geant4, the energy deposited in the detector is collected and turned into data for electron drift and scintillation photons. It is easy to misinterpret Geant4: Geant4 steps only when something interesting happens. For example, if a high energy muon enters the TPC and goes straight through, only ionizing the argon to leave a trail, Geant4 will take one big step throughout the detector and return the total energy deposited everywhere on the track. If you associate all that energy with one point, the interpretation will be incorrect. There are a couple ways to match the step length to the readout pitch: using a step limiter, dividing the total energy along the path, or using a segmented/voxelized geometry. Currently, LarSoft uses a voxelized geometry, but after our studies (p. 5-9), we concluded that it would be best to switch to using the step limiter instead.

After the tests, we came to the conclusion that Geant4 does a good job in carrying out the processes it was designed to do, although some follow-up studies need to be done on the shape of the electromagnetic showers. One must be careful about how one obtains data though, and must make sure to not misinterpret Geant4's results.

Hadronic Cross Sections:

Some current and future neutrino experiments might be interested in the cross sections for hadrons in liquid Argon. We ran tests with high statistics on Geant4, and plotted the results for pions, kaons, and nucleons. Here are the results:


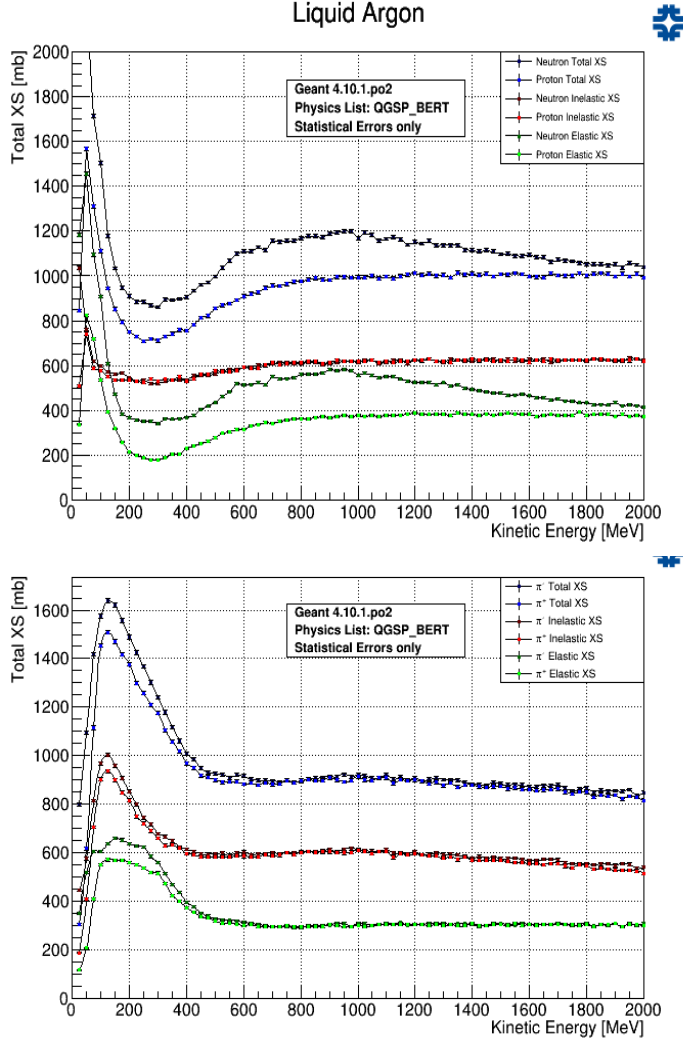


Figure 1. The elastic, inelastic, and total cross sections for pions (bottom), nucleons (top left), and kaons (top right) as a function of the kinetic energy. Since there is no experimental data to confirm these results, the integrity of these results might be based on the comparisons of Carbon data to Carbon simulations. See Figures 2-7.

Clear Delta resonances can be seen in the π^\pm and K-cross sections. In order to validate some of these tests, we did the same Cross Section tests on Carbon, and plotted experimental data on top of the graphs to see if Geant4 simulations matched experimental data (Figures 2-7). We used Carbon because there exists a lot of experimental data for Carbon. We found that the Geant4 tests did not match the experimental data for Kaons (see figures 4,5) on Carbon. The experimental data matched the simulated test fairly well for pions and protons (see figures 2,3, 6). We also ran tests for different physics lists: QGSP_BERT, QGSP_BIC, and QGSP_INCLXX, however there was no difference between the results of each list for this test.

The way we ran these tests was to turn on a beam of particles at the materials. We looked at a stepping action, and whenever the particle underwent an elastic process (hadElastic) or inelastic process (neutronInelastic, pi+Inelastic, etc.) we incremented the number of particles that reacted elastically and inelastically, respectively. Knowing the ratio of particles reacted to total particles for both inelastic and elastic interactions, we multiplied this ratio by a constant to give us the cross section:

$$XS = \frac{N_{interacted}}{N_{total}} \cdot \frac{A}{N_a \cdot d \cdot T} \quad (\text{eq. 1})$$

Where A is the atomic mass of element in the target, N_a is Avogadro's number, d and T are the density and thickness of the target. The total XS is simply the addition of the elastic and inelastic XS.




Figure 2. Comparison of Geant4 π^+ Carbon cross sections with experimental data. The results are consistent with each other. See bibliography for citations.




Figure 3. Comparison of Geant4 π^- Carbon cross sections with experimental data. The results are consistent with each other. See bibliography for citations.




Figure 4. Comparison of Geant4 K^+ Carbon cross sections with experimental data. The results aren't consistent with each other, suggesting one should take Geant4's treatment of Kaons with a grain of salt. See bibliography for citations.




Figure 5. Comparison of Geant4 K^- Carbon cross sections with experimental data. The results aren't consistent with each other, suggesting one should take Geant4's treatment of Kaons with a grain of salt. See bibliography for citations.




Figure 6. Comparison of Geant4 proton Carbon cross sections with experimental data. The few results we have are consistent with each other. See bibliography for citations.




Figure 7. We didn't find experimental data for neutrons on Carbon cross sections, but please enjoy the plot nevertheless.

Electromagnetic Showers:

We also looked at the shape of electromagnetic showers produced by 1 GeV electrons and photons in lead, iron, and liquid argon. We recorded both the radial energy profile and longitudinal profile to find the Moliere radius of lead, iron, and liquid argon, and the radiation length of liquid argon. For the Moliere radius, our results were all higher than the literature values.

	Simulated	Literature [cm]
Pb:	1.95	1.6
Fe:	2.79	1.72
lAr:	14.34	10.1

To find the Moliere Radius of Liquid Argon, we plotted (see figure 8) the total energy deposited in the detector as a function of radius from the central axis of the detector. We integrated the histogram (figure 8) until the partial integral was 90% of the total integral, and then recorded the radius at that point. The process was similar for lead and iron.

To find the radiation length for Argon, we looked at the longitudinal energy profile for many events. Figure 9 is a graph of the energy deposition as a function of depth in the detector. The result can be parameterized with the following equation:

$$dE/dt = E_0 t^a e^{-bt} \quad (\text{eq. 2})$$

Where $t = x/X_0$, a and b are free parameters, x is the depth into the detector, and X_0 is the radiation length. We fit this curve to the histogram for gammas, and the resulting radiation length was 10.9 cm. This is below the literature value of 14 cm for liquid argon.

The small difference in the shape of the gamma and electron showers is due to the ionization the electron does before it starts to shower. This deposits more energy in the beginning of the shower and leaves less energy left over for the tail of the shower.

This is not the first time somebody has tried to measure the Moliere radius of liquid Argon using Geant4 and come up with a higher answer. It is important to realize however that the literature values are based off of theoretical predictions, not solely data: it is impossible to make a homogenous calorimeter out of lead. Further studies need to be done to understand the discrepancy in the simulation values and the literature values.




Figure 8. A histogram of the energy of many showers deposited in the detector as a function of radius from the central axis. Notice the log scale of the y axis. We integrate this until we are at 90% of the total integral to find the Moliere radius, which was 14.34 cm.




Figure 9. This is a histogram of the energy of many showers deposited into the detector as a function of depth. It's shape is parameterised by (eq. 2), and from this longitudinal parameterization one can find the radiation length, which came out to be 10.9 cm. Only the histogram of the gammas has a good fit because the electrons leave some energy in the detector before they shower (ionization energy), throwing off the balance of the distribution.

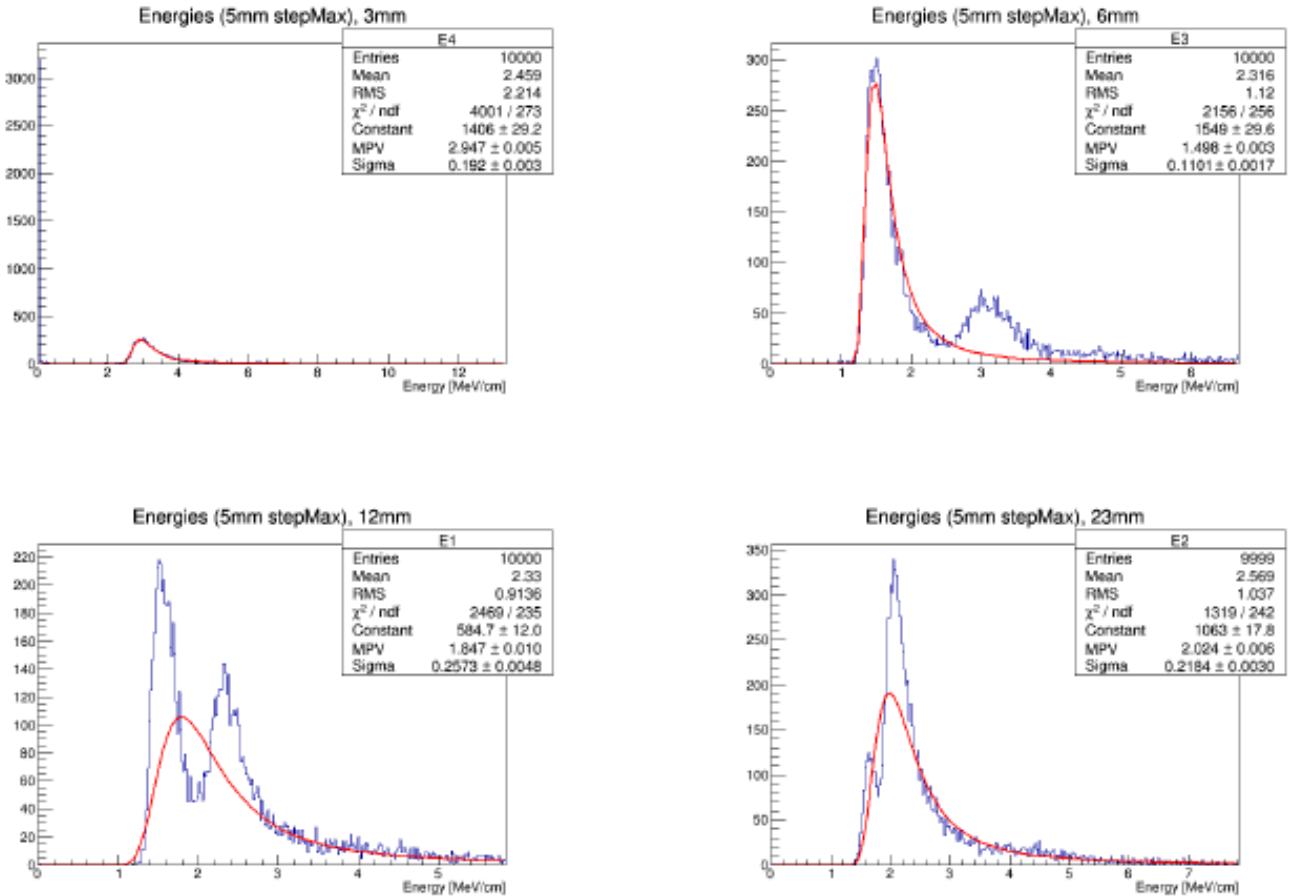
Step Limiter Comparison:

We did runs with 5GeV muons in liquid Argon, tracking the energy deposition in small slices of the detector. We expect landau curves because the muon approximates a minimum ionizing particle and according to Ereditato et al (see bibliography), the peak of the landau should be 2.1 MeV/cm. However, the muon in Geant4 takes steps that are bigger than the wire pitch of the TPC, which will lead to incorrect readouts. There are two options to deal with this problem: a step limiter or a readout geometry.

First, we used the step limiter, and the step limiter is set to a maximum step length of 5mm. Figure 10 measures the energy each muon deposited in the small 3mm slice of detector, and so on for slices of 6mm, 12mm, and 23mm. In the first plot, there is a

spike at 0 because a lot of the muons passed right through the 3mm slice without taking a step inside the slice. Therefore the energy deposited in that slice was never recorded, and so a lot of muons were recorded as leaving no energy in that slice of detector.

The rest of the plots in figure 10 show what look like double humped landau-ish distributions. In these slices, the step size is smaller than the slice width, so each muon will take a step in the detector slice and its energy will be recorded, so there is no peak at 0. However, many muons will take their last step in the slice but still deposit a considerable amount of unrecorded energy in that slice. This is why there are two peaks, and the first peak shrinks into the second peak as the slice width becomes much larger than the maximum step length.



We made the same plots for maximum step lengths of 1mm, .5mm, .3mm, and .1mm. The four slice widths are still 3mm, 6mm, 12mm, and 23mm.

Figure 10. The small slices of detector are made to represent TPC readouts. Since muons approximate minimum ionizing particles, we expect landau curves (the red fits) with peaks at 2.1 MeV/cm. When the step size is on the same scale as the wire pitch, the distributions are not what we expect.

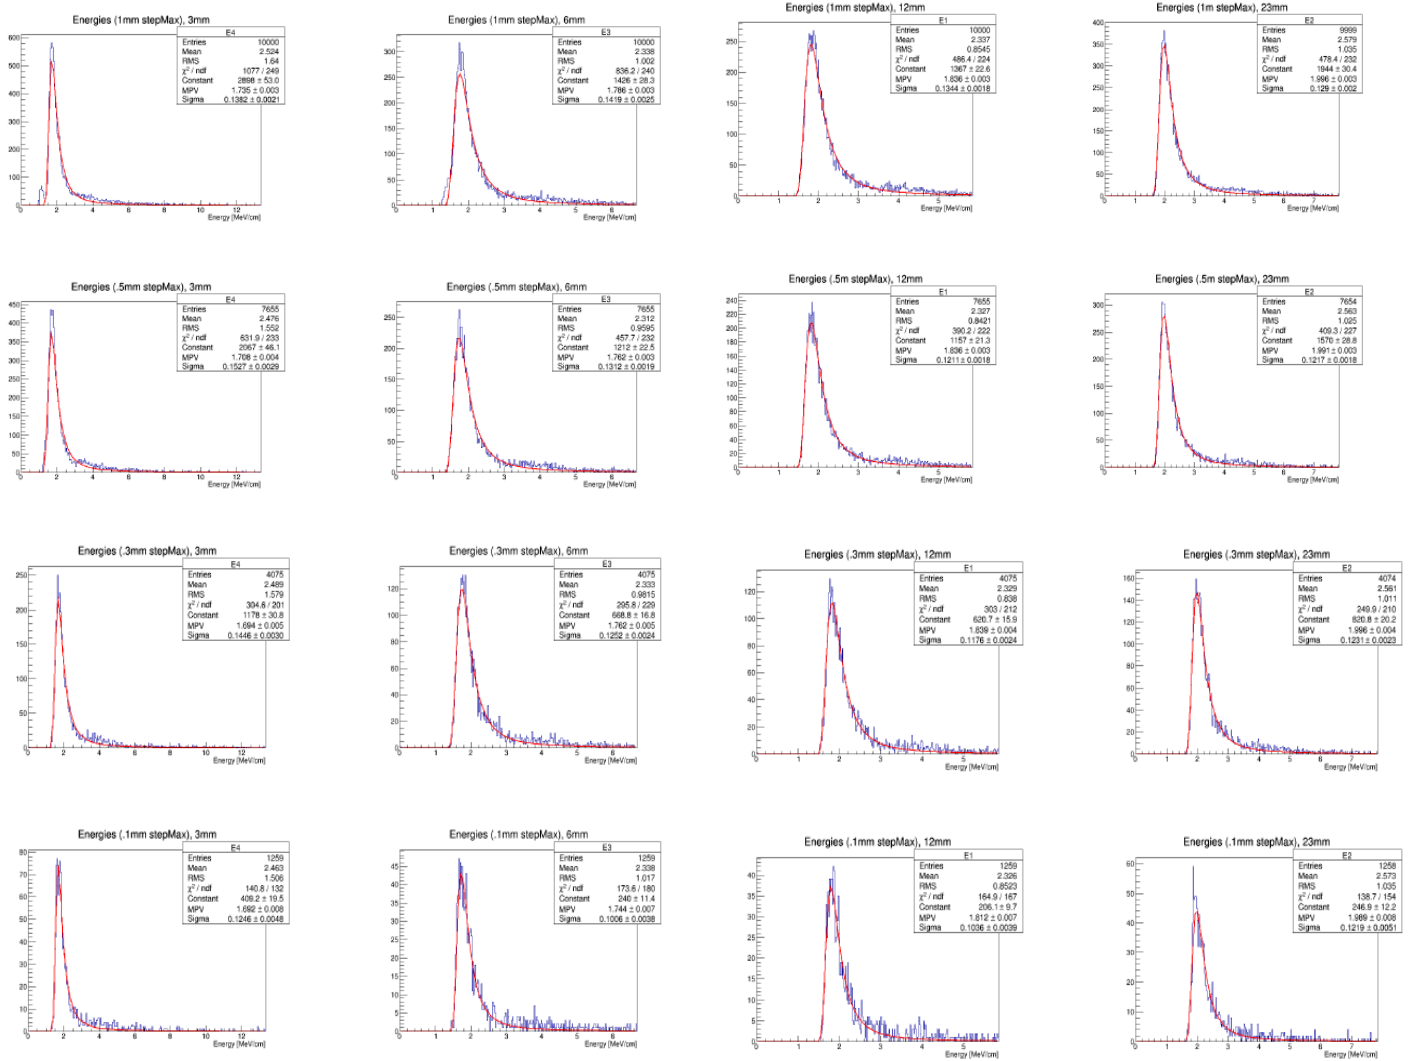


Figure 11. When the maximum step length is about a tenth the size of the detector slice, the resulting distribution is a landau distribution with a peak right where we would expect, at 2.1 MeV/cm.

The curve becomes a nice landau distribution when the maximum step length is about one tenth of the width of the slice of detector. So, when running G4 simulations, one should consider the max step length. However, decreasing the max step length increases the number of steps each particle has to take in the detector volume, so there can be a large computing price for having small step lengths. This is why the plots with smaller step lengths have less events than the others. Figure 12 is a graph of run time vs step size, (for the same number of events in each run) it is an inverse relationship (which is to be expected). Implementing a step limiter has no added memory costs. In conclusion, to simulate a TPC where the wire spacing is on the order of 5mm, a step size of .5 mm is sufficient for accurate physics simulation.




Figure 12. When you decrease the step size, each particle has to take more steps until it exits the detector. As a result, it takes longer to run for high statistics. The timing results for the Voxelized Geometry and divided hits layout are also on this plot.

Divided Hits Comparison:

Another option is to divide the large steps into smaller hits. This is like making the detector out of many smaller divisions, and when a large Geant4 step is made, the energy deposited for the step is distributed among the divisions that were passed. We did the exact same tests for the divided hits layout as we did for the step limiter, and the results were very different.

At the largest division size (figure 13), the energy distributions don't show the effects that the large step limiter shows, yet they do not match the landau distributions that we expect from a muon. However, as the division size decreases, as shown in figure 14, the distributions do not more closely resemble landau distributions. This is especially easy to see in the 6mm slice distribution (the second column), where a small peak does not go away near the base of the distribution. When the step limiter was down very low, all of the distributions closely resembled landau distributions, with good χ^2 values and a peak at 2.1 MeV/cm. This is not at all true for the smallest division sizes.

The reason the distributions do not more closely resemble landau distributions is that by dividing the total energy taken over the step, we are getting rid of all of the small fluctuations that make up the tail of the landau. This keeps a landau curve from fitting nicely to the shape of the distribution, even when the division size is very small.

Dividing the steps up into smaller hits takes up a lot of memory (see Figure 15). At large division sizes, Geant4 doesn't take up very much memory, but when the divisions get small, Geant4 takes up a huge amount of memory. This is very dangerous; especially if you want to run a simulation job on the grid, using so much memory can get very expensive. And as shown by figure 14, despite using more memory and decreasing the division size, the simulation does not do a better job of representing the actual physics in the detector. All taken together, using the step limiter seems to be a better choice to get the physics right in Geant4.

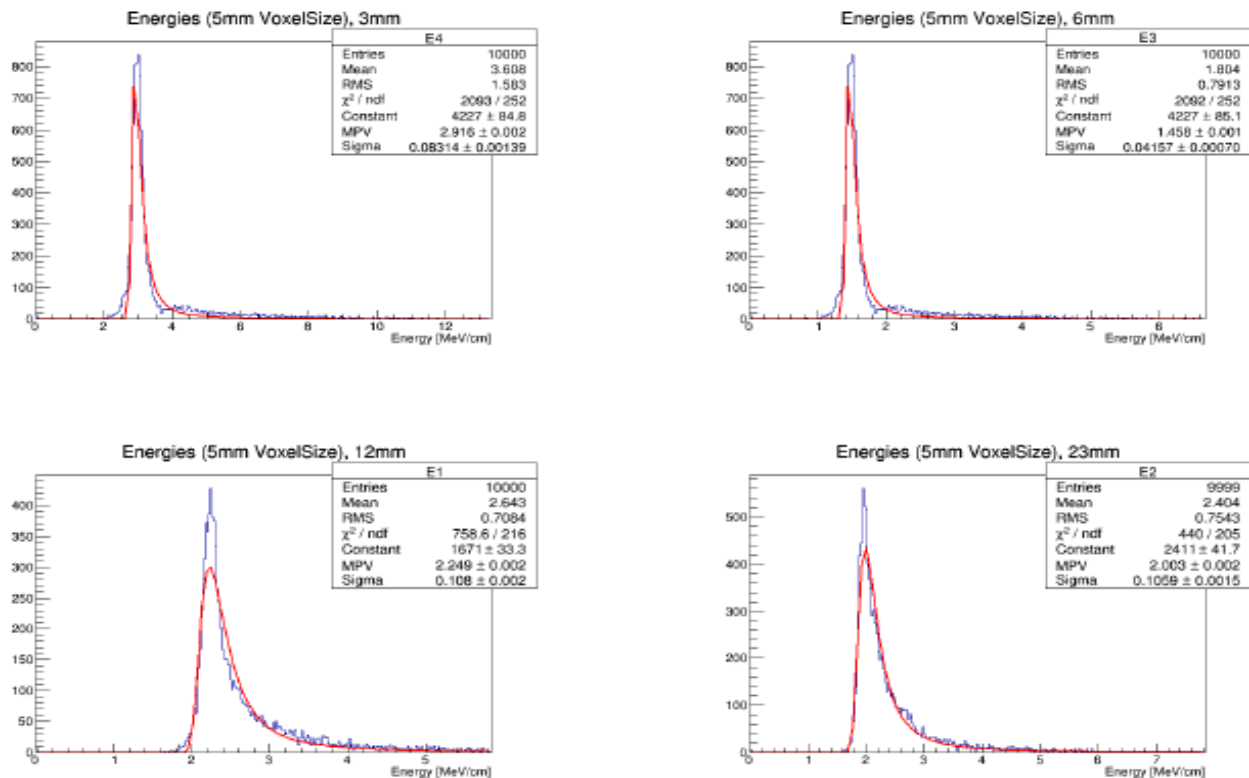


Figure 13. We do not expect the largest division size to give us great readouts, and they don't, however as we decreased the division size the readouts stayed the same.

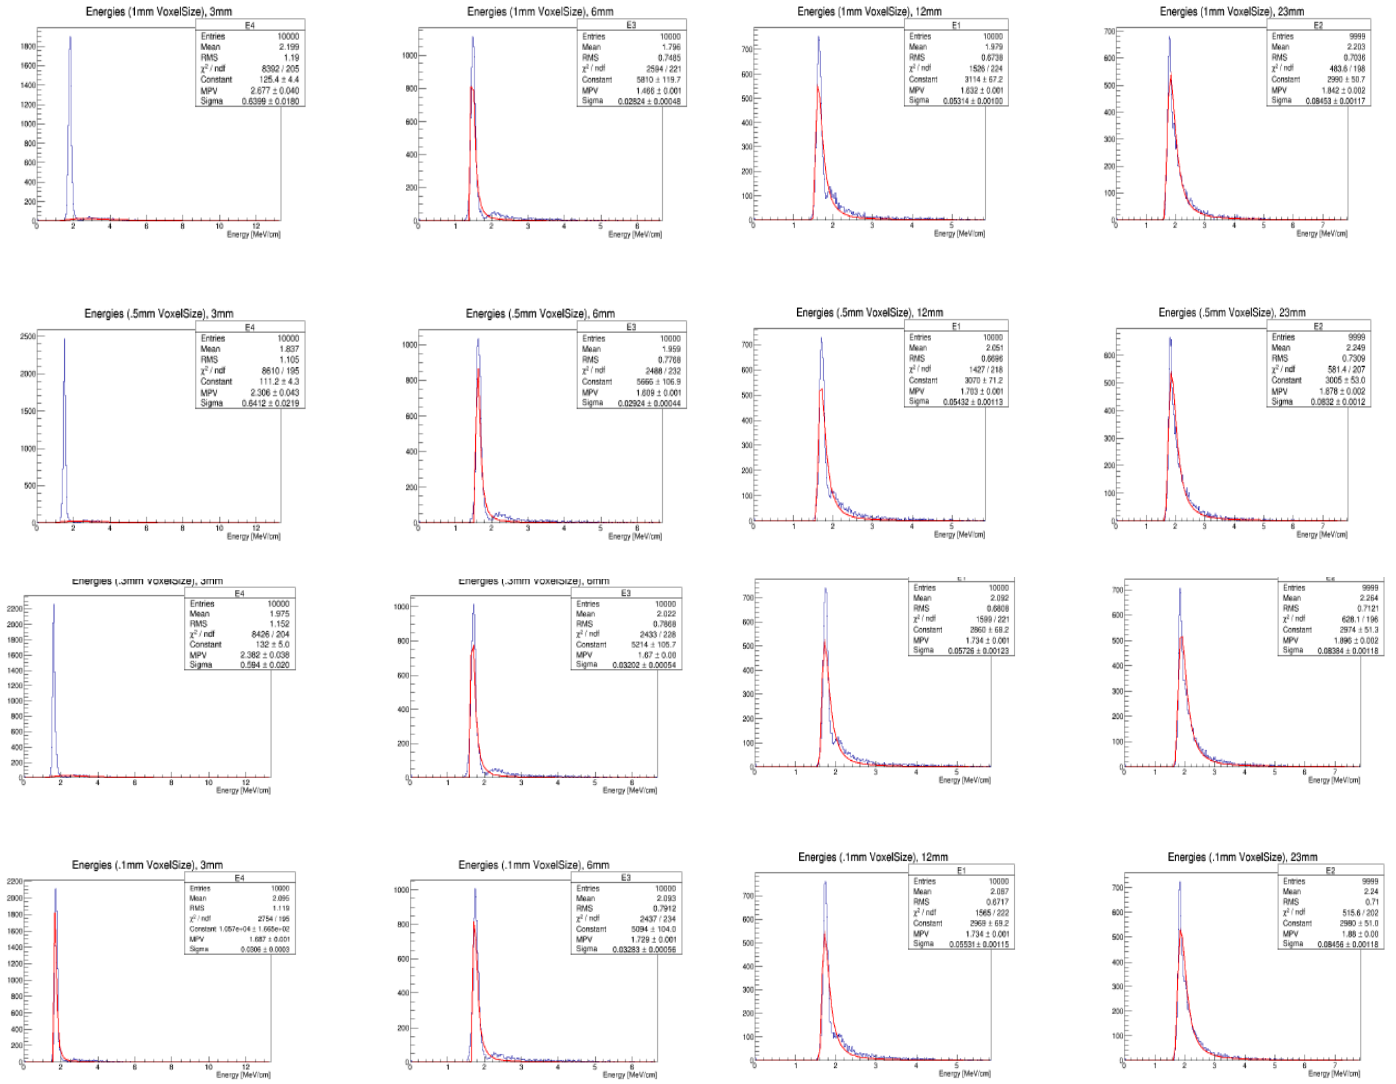


Figure 14. There seems to be no improvement in the energy distributions across a small slice of detector when the division size is decreased.

As the memory cost gets quite large for smaller divisions, the time it takes to run does not get very large. The time to run with each division size was under .06 sec/event, all the way down to the smallest division. This is extremely fast compared to the step limiter, however the step limiter uses very little memory in comparison. So there is a trade-off there, however the results from the landau fits of the Step-limiter distributions combined with the enormous memory cost of the divided hits layout makes an excellent case for using a step limiter instead of a divided hits layout.




Figure 15. Although the divided hits layout ran very quickly (see figure 12), it used up a lot of energy when the division size got small. For comparison, the step limiter used no additional memory.

Voxelized Geometry Comparison:

The final option to match the step length to the readout pitch is to use a segmented, or voxelized, geometry. This involves making the detector out of little cubes. Geant4 is forced to step at each boundary, and so using this geometry is similar to using a step limiter. Therefore, the results are not much different than the step limiter results above (see figure 11). However, if you match the voxels to the readout pitch, the result is a nice landau distribution (see figure 16).

Unfortunately, the wire planes of most liquid argon TPC detectors are not so simple. The planes are normally not at right angles to each other, and often there are three wire planes, so it is impossible to perfectly match the cubic voxels to the readouts planes. Therefore, the readout geometry acts effectively as a step limiter, however with some extra costs. It takes longer to run with the voxelized geometry (.9 sec/event with .1 mm voxels) than with the step limiter (see figure 12). It costs only about 10MB of additional memory to run, so it is nowhere near the memory costs of the divided step method; it is about the same memory cost as the step limiter.

Currently, LarSoft uses a voxelized readout geometry like the one I have described to match the step length to the readout pitch. Since the step limiter is quite simple to implement, and the segmented geometry takes considerably longer to run, it might be beneficial to implement a step limiter in the LarSoft Code and remove the voxelized geometry.




Figure 16. When the voxels are matched to the readout slices, the distribution of energy depositions closely resembles the landau curve we expect. Above is a 5mm voxel fit exactly to a 5mm readout slice, and the result is very nice. However, there are multiple wire readout planes in most liquid Argon TPC detectors, all at different angles (not 90 degrees), making it impossible to perfectly match the voxels to the readout planes. The rest of the segmented geometry data is not shown because the segmented geometry otherwise acts as a step limiter.

Muon Separation in lAr (50 MeV):

One of the difficulties in detector physics is differentiating between oppositely charged particles, especially in the absence of a strong magnetic or electric field. We ran some studies to figure out how to separate a μ^- from a μ^+ . First of all, when going through the detector, the muons go through different processes, because the detector is made of matter, and antimatter and matter interact differently in the material.

Processes	μ^-	μ^+
Decay*	$e^-, \bar{\nu}_e, \nu_\mu$	$e^+, \nu_e, \bar{\nu}_\mu$
Lifetime	.58 μ s	2.2 μ s
Probability	25%	100%
Capture	ν_μ, γ^{**}	none
Probability	75%	

*For μ^- , this is decay in orbit

**There can be more than one gamma

When the μ^+ enters the target, it acts like a minimum ionizing particle and ionizes the Argon. It gets stopped and decays into a positron and two neutrinos. As soon as these three particles are created in the simulation, Geant4 puts them on the stack. We used a Stacking Action to find the Kinetic energy of the e^+ and the time just as it is created, and put these values into histograms.

The Kinetic Energies (figure 18) of the positron are what we expect. The muon has a rest mass of 106 MeV, and that energy has to be split among the neutrinos and the positron, conserving momentum. Since the neutrinos have such little mass, the positron can't have too much energy or else momentum conservation wouldn't be possible. This is why the kinetic energies of the positrons falls off at around 50 MeV. The timing (figure 17, in ns) of the positron is also what we expect. It represents when the muon decays. The histogram fits a very nice exponential decay, the standard for decay rates, and has a mean of 2.2 μ s, which is the literature value for the mean lifetime of a muon.

The μ^- decay is a little different. First of all, every negative muon that entered the target was captured by an Argon atom. There, it was in "orbit" where it either decayed or was captured by the nucleus. 25% of the time it decayed. Similarly for the positive muon's positron, we measured the timing and energy of the e^- that came from the decay of the negative muon. The Kinetic Energies (figure 20, in MeV) are as expected, similar to the positive muon, and the difference in the shapes of the distributions is due to the




Figure 17. This is the time (in ns) from the beginning of the event that the positron from muon decay was put on the stack. This coincides with the timing of the muon decay. The equation above is the classic decay equation, where t_{μ^+} is the mean lifetime of the μ^+ in liquid Ar.




Figure 18. This is the kinetic energy (in MeV) of the positron just as it is created. Since two neutrinos have to have the leftover energy the positron takes and momentum has to be conserved (the muon decays at rest), the positron cannot have too much of the energy or else momentum will not be conserved.

decay in orbit. The Argon nucleus can have some of the kinetic energy released in the muon decay, which slightly changes the shape of the distribution. The times (figure 19, in ns) follow the expected decay distribution but the mean is 0.58 μ s instead of 2.2 μ s. This is because the negative muon is matter, and so it has weak interactions with the matter around it, shortening its lifetime. Suzuki et al. measured the lifetime of the muon in liquid Argon, getting 0.54 μ s, with which our value of 0.58 μ s is consistent.

The negative muon did not always decay. It was always captured by the atom and put into “orbit.” It had two different possible processes: decay in orbit or nuclear capture. 75% of the time the muon interacted with the nucleus and created a muon neutrino and gammas, turning a proton into a neutron, and possibly expelling some protons and neutrons in the process. We looked at the gammas coming out to see the timing of when the events happened and the energy of the gammas. The timing was the same as the timing for decay in orbit. The kinetic energy for the gammas can be seen in figure 22.

We also made a plot to show the different atoms and isotopes that resulted from the muon capture (in figure 21). The graph represents the relative probability of the muon creating one atom from Argon over another. Cl40 is the most abundant, and therefore the most likely. This is when just a proton becomes a neutron. However, a large variety of elements and isotopes can be produced from the energetic nuclear capture.




Figure 21. Whenever there was nuclear capture of a muon, there was a chance for nucleons to be kicked out of the nucleus. Geant4 always put the resulting nucleus on the stack, where we were able to tell what it was. So, we put the different nuclei into a histogram to see what the relative likelihood of creating other nuclei was. A lot of the time Cl40 was formed, so the proton-neutron conversion was all that happened, but most of the time the process kicked out a few neutrons and protons too.




Figure 19. This is the time (in ns) from the beginning of the event that the electron from muon decay was put on the stack. This coincides with the timing of the muon decay. The equation above is the classic decay equation, where t_{μ^-} is the mean lifetime of the μ^- in liquid Ar.




Figure 20. This is the kinetic energy (in MeV) of the positron just as it is created. Since two neutrinos have to have the leftover energy the positron takes and momentum has to be conserved (the muon decays at rest), the positron cannot have too much of the energy or else momentum will not be conserved.




Figure 22. This is the Kinetic energy (in MeV) of gammas emitted from nuclear capture. We didn't really expect a specific pattern, but this gives a good general idea of the energies of these gammas.

When the muon was captured by the atom and put into orbit, there was a sort of “flash” of some gammas and some electrons. We looked at the energy spectrum of the gammas coming out (figure 23). With the isolated peaks, it looks like spectral lines. This makes perfect sense, because the muon's electromagnetic interaction with the atom should mimic an electron's electromagnetic interaction with the atom, and that would excite the atom which would emit light with the spectral energies. The timing for these gammas was much before nuclear capture or decay, on the order of less than a nanosecond, right when the muon enters the “orbit” of the atom.

The goal is to be able to look at detector tracks and know whether you are looking at a positive or negative muon. Since all the μ^- were captured by the atom, emitting some light and none of the μ^+ did that, if the detector could detect that light then that would be a great way to determine the charge of the muon. However, most of this light has very low energy, with a mean of .16 MeV, and might be very hard to detect with your detector. If they both decay, the μ^- will decay into an electron, while the μ^+ will decay into a positron which will annihilate. If you can find the annihilation, then you know it was a μ^+ . We plotted the energies of the gammas from e^+ annihilation, and the energies of the resulting gammas is extremely consistent, at .52 MeV which is consistent with the literature value for the rest mass of an electron, .51 MeV (an electron and a positron decay into two gammas, so each gamma should have .51 MeV). This low energy level also may be difficult to detect, but it is a very consistent value. To see if the muon was captured by the nucleus, you could look for the relatively high energy gamma (figure 22), which could pair produce and maybe shower later in the detector. Or you could look for the absence of the michel electron. Then it would be a μ^- .




Figure 23. This is the energy distribution of the gammas emitted when the muon was captured by the atom. It resembles a spectral plot, which is exactly what it is. It would be a useful tool to identify when a muon gets captured by an atom, however the energy of the photons might be too low to be detectable (with a mean of .16 MeV).




Figure 24. This is the energy of the gammas from the positron electron annihilation. It is very consistent, so it might be useful for detecting positrons from muon decays, however the signal might be too weak to be detected.

Pions (50 MeV) in lAr:

We also looked at π^+ and π^- in liquid Argon to see what processes they underwent and what we could do to separate them by charge in an experiment, and separate pions from muons. The table below shows the different pion processes Geant4 does and the percent of times they happen.

Processes	π^-	π^+
Decay	$\mu^-, \bar{\nu}_\mu$	μ^+, ν_μ
Lifetime	?*	26 ns **
Probability	2%	95%
Inelastic Probability	9%	5%
Capture Probability	89%	none

*For π^- , there were not enough statistics to determine the lifetime

**The measured π^+ lifetime is consistent with the literature value of 26 ns.

The pions decay into muons (although decay into an electron is possible but exceedingly rare) and we have a pretty good understanding of how muons interact in liquid Argon (see above). However, the muon might leave the detector before it gets a chance to decay. However, if you see a pion decay, it is likely to be a positive pion because they decay much more often than negative pions. However, to be sure it was a positive pion, you could look for the positron from the resulting μ^+ decay or the two gammas from the e^+e^- annihilation. If there is a muon capture, then the decay might have been from a negative pion.

The difference in the Inelastic process for the pions is that π^- will convert a proton into a neutron, while π^+ will convert a neutron into a proton. It doesn't make sense to distinguish between these two, so telling the inelastic processes apart from each other is rather difficult. In both inelastic processes and the Nuclear Capture of the π^- , lots of junk (protons, neutrons, deuterons, etc.) is expelled from the nucleus, just like the capture of μ^- .

Therefore, if a particle comes into your liquid Argon TPC and creates a sort of hadronic shower, with neutrons and protons and gammas, it would be very difficult to tell if it was a μ^- or a pion. If a particle comes into your detector and decays, there are plenty

of signs to tell you what it was. If soon after the primary decay, there is a track and a secondary decay, then you know it was a pion decaying into a muon. By further inspecting the Michel electron of the secondary decay and seeing if it annihilates or not, you could distinguish between π^- and π^+ . If soon after the primary decay there is what appears to be a small Hadronic shower, then a μ^- created that shower and therefore the primary particle was a π^- . And, if there is a decay but no secondary decay or shower, then the primary particle might have been a muon, see above for distinguishing between muons. You can also use timing to distinguish between muons and pions because pions have a much shorter lifetime than muons.




Figure 25. This is the time (in ns) from the beginning of the event that the muon from pion decay was put on the stack. This coincides with the timing of the muon decay. The equation above is the classic decay equation, where t_{π^+} is the mean lifetime of the π^+ in liquid Ar.

Bibliography

- Allardyce, *et al.* "Pion reaction cross-sections and nuclear sizes." Nucl.Phys.A209 (1973). p. 1-51.
- Ashery *et al.* "True absorption and scattering of pions on nuclei." Phys.Rev.C23 (1981).
- Bugg, *et al.* "Kaon-Nucleon Total Cross Sections from 0.6 to 2.65 GeV/c." Phys.Rev.168 (1968). p. 1466-1475.
- Clough *et al.* "Pion-Nucleus Total Cross Sections from 88-MeV to 860-MeV." Nucl.Phys.B76 (1974). p. 15-28.
- Ereditato, A., and A. Rubbia. "Ideas for a next Generation Liquid Argon TPC Detector for Neutrino Physics and Nucleon Decay Searches." 27 Apr. 2004. Web.
- Friedman *et al.* "K⁺ nucleus reaction and total cross-sections: New analysis of transmission experiments." Phys.Rev.C55 (1997) p. 1304-1311.
- Gelderloos *et al.* "Reaction and total cross-sections for 400-MeV 500-MeV pi⁻ on nuclei." Phys.Rev.C62 (2000).
- Schwaller *et al.* "Proton Total Cross Sections on ¹H, ²H, ⁴He, ⁹Be, C and O in the Energy Range 180 to 560 MeV." Nucl.Phys.A316 (1979). p. 317-344.
- Suzuki *et al.* "Total nuclear capture rates for negative muons." Phys.Rev.C35 (1987).
- Wilkin *et al.* "A comparison of pi⁺ and pi⁻ total cross-sections of light nuclei near the 3-3 resonance." Nucl.Phys.B62 (1973). p. 61-85.